

# Elasticsearch Cloud Prototyping




We identified a few potential risk and opportunity areas when investigating moving CMR to the Cloud. One of those areas was transitioning Elasticsearch to the cloud. We looked into one potential opportunity which was to use Amazon's Elasticsearch service instead of managing our own EC2 instances. We ruled that out based on the limitations listed here: <http://docs.aws.amazon.com/elasticsearch-service/latest/developerguide/aes-limits.html>. Namely the lack of support for custom plugins, a maximum cluster size of 10 nodes, and the version of Elasticsearch is well behind what we use today. We also identified a potential risk being the underlying data storage used for the cluster.

## Elasticsearch Storage

There are 3 different storage options that could be used for the Elasticsearch cluster.

1. Local SSD
2. EBS General SSD
3. EBS Provisioned IOPS

Each storage option has its own set of advantages and disadvantages.

Storage Type	Permanent	Performance	Number of Replicas	Relative Cost	Misc Notes
Local SSD		Roughly equivalent to current performance	2*	1	The main issue is that when a node is taken down the storage is gone. If we were to lose two nodes simultaneously we would most likely lose both the primary and replica for several shards. Then we would need to recover the data from a snapshot and then reindex any data since that snapshot. As a mitigation we could choose to have 2 replicas which would effectively increase the number of EC2 instances required by 50%. If we were to use 2 replicas the cost is just about the same as using EBS General SSD storage.
EBS General SSD		Likely slower than current performance.	1	1	We would need to test performance to ensure it is adequate. Advantage is that data will not be lost if a node goes down. Need to research how we would handle an instance dying to ensure the EBS Volume is re-assigned when an instance is started back up in place of the failed node. Similarly would need to determine how to handle EBS volumes as part of patching since NGAP uses immutable infrastructure.
EBS Provisioned IOPS		Would need to test. Latency will be worse than local SSDs, but high IOPS may make up for the worse latency.	1	1.27 (1000 IOPS) 1.43 (2000 IOPS)	27-43% higher cost than other storage solutions if configuring between 1000 and 2000 IOPS. Will provide better performance than EBS General SSD storage, but not sure how it compares to local SSD due to latency. Same advantages and disadvantages as EBS General SSD.

## Test Plans

We determined that we want to run a series of tests using General Purpose SSDs as well as Local ephemeral SSDs.

1. Snapshot and Recovery
2. Rolling restart of cluster
3. 24 hour search load test

More details for the tests can be found in the individual tickets within the epic

 [CMR-2601](#) - JIRA project doesn't exist or you don't have permission to view it.